ON THE ROAD TO THE INTERNET

Dr. Vinton G. Cerf
Vice President and Chief Internet Evangelist, Google Inc.

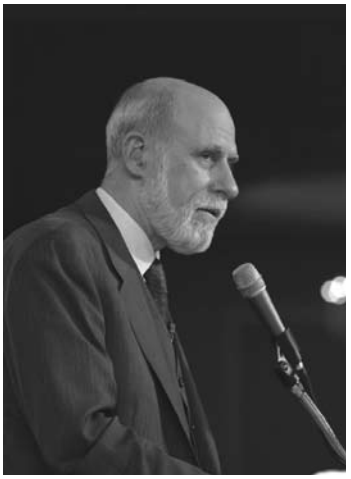

Members of the Japan Prize Foundation, distinguished guests, ladies and gentlemen, it is an honor to have this opportunity to share with you my personal journey that led to my participation in the invention and development of the Internet. The organizers of this event have asked me to outline some of my memories of childhood, academic experiences and later work that led up to the creation of the Internet. It is also important to appreciate the long and difficult road after the basic invention to its present state as a global information infrastructure.

As a child growing up in the 1940s and 1950s, I was always surrounded by books and by a family that believed strongly in education. I was encouraged to work hard at school and to do my best for my teachers. While my mother and father were college graduates, they did not have advanced degrees, but they recognized the importance of education in their own lives and reinforced this belief in their children. I was the oldest of three brothers but until the age of five, I was an only child. I spent a lot of time with my mother during that time and she imbued me with a sense of humor and an appreciation for music. A favorite pastime was to listen to classical works played on the radio and to try to guess the name of the piece and the composer before the announcer reminded the audience of the origin of the piece just played. We also spent time at museums and, like many children, I was fascinated by dinosaurs, Egyptian mummies and other things from the ancient past. I was an inveterate reader and accumulated a respectable library of books as a child. I recall reading books like One, Two, Three, Infinity by George Gamow, and Microbe Hunters by Paul deKruif. A favorite book around age 12 was The Boy Scientist by John Lewellen. It was a practical book that described experiments and explained various scientific concepts that I found fully absorbing. By that time, I had acquired a chemistry set (this would have been around 1955 when American chemistry sets had quite an extensive array of chemical materials). I would spend hours trying out various mixtures to see what would precipitate out. Of course, like many boys my age, I was fascinated by pyrotechnics and delighted in making plaster of Paris volcanoes that could be set off using hypergolic materials, sulphur, powdered aluminum and powdered magnesium. Small rockets made by filling empty rifle bullet shells with match heads were another dangerous but equally fascinating pastime. As I look back on those days, I am amazed that I and my inquisitive young friends did not damage property or ourselves with some of these experiments.

Fascination with chemistry was matched by a deep interest in mechanical constructions using an Erector set. To this was added an abiding interest in mathematics. I recall complaining to a 5th grade teacher at age 11 that the arithmetic I was being taught was boring. I received a 7th grade algebra book in return and had a marvelous summer solving every problem in the

On the Road to the Internet

Vint Cerf
Japan Prize 2008

book. I particularly liked the word problems because they seemed like little mystery stories. You had to figure out just what X was at the end.

My interest in mathematics led quickly to an interest in computers. By good fortune, my father had a good friend who was writing software for a project of the US Government called SAGE for Semi-Automated Ground Environment. This tube-based computer system accepted input signals from radars located in northern Canada placed at what was called the Distant Early Warning (DEW) line. The idea was to detect automatically any attempts by the Russians to fly bombers over the North Pole to attack the US or Canada. The computer that was used to do this work was so large that it filled several rooms. One actually walked inside the computer to use it. My first encounter with this system took place in 1958 at a placed called the System Development Center in Santa Monica, California. Two years later, my best friend in high school, Steve Crocker, got permission to use a computer called a Bendix G-15 located at the University of California, Los Angeles. We were permitted to use the machine when it was not in use by others so we often spent nights or weekends preparing programs for the computer to run. Usually we were interested in plotting solutions to transcendental equations that could not be solved easily in closed form but could be plotted on paper using the values produced by the computer program.

During this high school period in my education, I took every academic advanced placement or enrichment course I could and participated in the math club contests and things like the Knowledge Bowl in which high schools competed with each other. It was a thrill to be part of the team that won these events for the honor of the school. I also served in high school as the editor of the creative writing magazine called the Winged Pen. From this experience I took away an interest in creative writing and poetry that has stayed with me to the present. I think it would be hard to over emphasize the effect of my early school years on my interest in science, mathematics, literature and history. The gift of books and reading has borne dividends for decades and continues to be a source of great pleasure for me.

I often envy young people today who are exposed to computers and the Internet at such early ages. Even six year olds seem to find things to do with computers these days. I was seventeen before I got to program a computer and of course, had to participate in inventing the Internet before I could use it!

My father also believed that learning a second language was an important experience so he engaged a tutor from Germany when I was still in Junior High (about 8th grade, at age 13). He would come every Wednesday evening and we would spend two hours reading and conversing in German after which I would have to recite something in that language for my parents and then we would have dessert.

It was my good fortune that my father worked at a company then called North American Aviation. It had a scholarship program and I was lucky to win a full four year scholarship that allowed me to take advantage of my acceptance at Stanford University. North American had a number of subsidiaries and I was able to work at many of them during summers as a high school and college student. One subsidiary, Atomics International, designed nuclear power

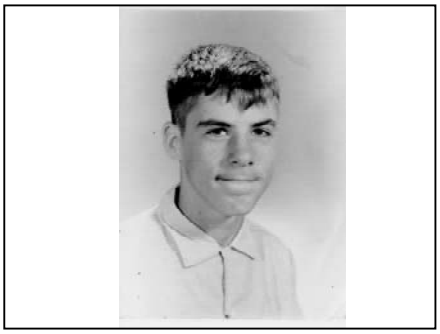

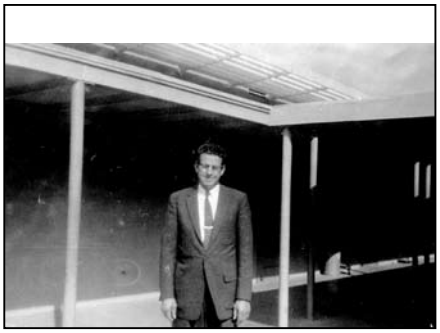

systems, and I worked there during one summer as a high school student. Another subsidiary, Rocketdyne, was deeply involved in the American space program, notably the Apollo effort. I had a small role to play as a recently-graduated high school student during the summer of 1961 analyzing the test data from the huge F1 Apollo Saturn V rocket system. These massive liquid fuel engines produced 1.5 million pounds of thrust each. They were test-fired in the Santa Susanna Mountains north of Los Angeles and I worked on the analysis of the data to try to determine whether these engines would survive the short but critical boost phase of an Apollo launch.

By the time I began my undergraduate work at Stanford University in 1961, I had a strong interest in mathematics and computing, and took a curriculum heavy with these subjects. Stanford is a liberal arts school and its curriculum required what was then called the History of Western Civilization. I had to read many, many books from Greek and Roman times up through the European Renaissance and the so-called Age of Reason. To this day, I am grateful that Stanford required me to read these books as it seems unlikely that I might have chosen to read these later in life. Stanford also had a foreign campus program and I chose to go to Germany in my sophomore year. Exposure to a culture and language different than my native American expanded my world view noticeably. We had classes in German and took many field trips to cities around Europe. We learned about history, geography, architecture, literature and even linguistics during this six month period. I was fascinated with the evolution of Old High German and Old English into their modern equivalents today.

During the summers, I would work at various of the subsidiaries of North American Aviation, including its Space and Information Systems Division where I wrote computer programs in support of the Apollo space program and again at Rocketdyne.

After my undergraduate education in mathematics, I decided I wanted to get some practical experience with computing so I applied to work at IBM in Los Angeles and was accepted into their systems engineering program. I wound up working in the Los Angeles Data Center running their Quiktran time sharing system in 1965. This was very early in the history of time-sharing systems that had been invented only a few years before at the Massachusetts Institute of Technology by a professor, John McCarthy and his colleagues. McCarthy later went to Stanford where I encountered him when I was an undergraduate there.

By good fortune, I met my wife-to-be, Sigrid Thorstenberg, in late 1965. We had the same hearing aid dealer and he introduced us one Saturday morning. We became very good friends very quickly, discovering many shared interests and by September 1966, we were married in the Wayfarer's Chapel in Palos Verdes, Calfornia. Sigrid was an interior designer and this creativity colored all the places that we have subsequently lived and without doubt had a great impact on the genetic heritage of our two sons.

After two years with IBM where I learned a great deal about practical computer systems, especially operating systems, I felt a strong need to return to school for advanced training in computing. My good friend, Steve Crocker, once again helped me by introducing me to his advisor, Prof. Gerald Estrin, who was an enthusiastic mentor and encouraged my return to

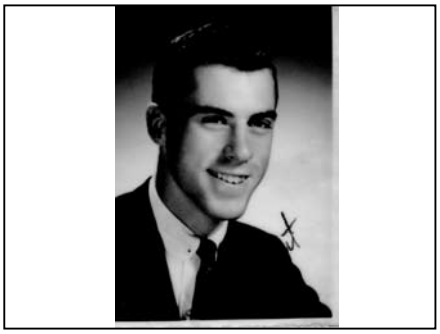


NAME
COME

academic life.

During this period, I took many formal classes in computing, queueing theory, programming languages and operating systems. My dissertation work took me into parallel programming and the transformation of serial programs into parallel ones suitable for execution on multi-processor architectures. Ironically, the most recent computer chips that are being produced have multiple processors on each chip, making this early work suddenly relevant again.

By 1968, however, the US Defense Advanced Research Projects Agency (DARPA) had decided to pursue the design and implementation of a packet switched computer network to encourage resource sharing among the many computer science departments that DARPA funded to do research. UCLA was to be the first site of the ARPANET and I began working with my colleagues, including my good friend, Steve Crocker, on software and protocols that would be suitable for allowing computers on the network to exchange information and support remote time-shared access. During the period from 1969 to 1972, a small team of graduate students worked under the direction of Prof. Leonard Kleinrock, who was the principal investigator of the Network Measurement Center. My job was to write the software that allowed the UCLA computer to inject patterns of traffic into the network and to capture the performance results which would then be compared to the mathematical models that Prof. Kleinrock and his other students had developed. In parallel with the performance work, Steve Crocker led the effort to develop the communication protocols that would be used to support inter-computer applications.

It was in early 1970 that I met Robert Kahn who was one of the key architects of the ARPANET's packet switches, called Interface Message Processors or IMPs. Bob was then working for Bolt Beranek and Newman (BB&N) in Cambridge, Massachusetts. He came to UCLA to run tests on the new network and to demonstrate that under certain conditions the network could "lock up." I helped him in this work by writing software to generate artificial traffic and inject it into the network while extracting from the network detailed performance information. This early and fruitful collaboration foreshadowed our later work together on the Internet.

By 1971, protocols had been developed to allow time-shared access to remote computers on the network, file transfers and electronic mail. It is hard to believe that before that time, email did not exist except for users of the same computer who could leave files for each other on the machine they shared. By 1972, DARPA sponsored a public demonstration of the ARPANET at the International Conference on Computer Communication (ICCC) held in Washington, DC in October 1972. The director of the DARPA Information Processing Techniques Office, Larry Roberts, asked Robert Kahn to lead this demonstration effort and a great many graduate students from the two dozen or so universities on the ARPANET participated in preparing public demonstrations of applications that could be supported on the network.

About that time, I left UCLA to join the Stanford faculty with a joint appointment in Computer Science and Electrical Engineering. My friend and later partner, Bob Kahn, left Bolt Beranek and Newman to join DARPA and Larry Roberts left DARPA to join BBN and to


KINGDOM

become the CEO of Telenet, one of the first commercial packet switching networks.

The invention and subsequent evolution of the Internet rests on a foundation that reaches into the early 1960s and by some measures even further back, arguably as far as the mid-1900s with the invention of the Telegraph. Indeed, a popular historical book by Tom Standage entitled The Victorian Internet is, in fact, all about the telegraph and its rapid adoption and global reach. Following that is, of course, the invention of the telephone and then radio. Every one of these inventions and the technologies associated with them has had some influence on the emergence of the Internet.

It is a truism that inventions can only happen when conditions are right for their successful implementation. These conditions may be technical, economic, social or political or a mixture of the three. In some respects we can attribute the conditions leading to the Internet to all four of these categories.

Once at Stanford, I split my time between teaching classes and initiating research on a new network design that led to the Internet. The origin of that work came approximately in the Spring of 1973 when Bob Kahn came to visit from DARPA and told me about the packet radio and packet satellite network plans he was pursuing. The idea was to demonstrate that packet switching could work in a variety of media. About this same time, Robert Metcalfe who was then at XEROX Palo Alto Research Center (PARC) was inventing the Ethernet. My graduate students were aware of this work because some of them had jobs at PARC and some PARC staff were doing advanced degree work at Stanford. Consequently, as Bob and I began to explore ideas for the design of the Internet (as it was later called), we had the benefit of knowing about other network research in addition to the work that DARPA was funding.

Xerox PARC had brought a number of people from SRI International which has been a pioneer in online computer use through a project led by Douglas Engelbart. The so-called Augmentation Research Center was exploring ways to augment human capability by developing significant computer tools for knowledge processing. They ALTO workstation. They connected over 250 of these using their Ethernet and they built a wide range of applications for these machines. These developments were very much in Bob Kahn's mind and mine as we discussed his ideas for open networking in mobile and fixed environments including ships at sea, mobile vehicles on the ground and on planes in the air. It was in this context that the design of the Internet evolved.

1973 witnessed the birth of many things including the birth of our first son, David Cerf, in September of that year.

Bob Kahn and I also took advantage of colleagues in France, particularly Louis Pouzin, who was working on the design of a packet switched network in 1973-1974. In fact, one of his team, Gerard LeLann, was to come to Stanford to work with me and my team during 1974. Bob Kahn and I wrote a paper about the design of the Internet in late 1973 which was published in May 1974. We talked about an "internetwork" at that time. During 1974, I led a team of graduate students and visiting colleagues in the detailed design of what we called the Transmission Control Protocol (TCP) and in our final document released in December 1974,

we described the details of how this end-to-end computer protocol would work in the network of networks we came to call the Internet.

During this same time period, the ARPANET continued to thrive and I was able to demonstrate its power remotely from places like South Africa and Brazil using dial-up modems and the telephone system to connect to suitably-equipped ARPANET nodes.

Once the TCP specifications were done, we began implementation early in 1975 with colleagues at BB&N and at University College London. We quickly discovered mistakes that we had made in the protocol design and proceeded to amend these, including the addition of the so-called "three way handshake" to assure that both ends of a TCP connection would know how to sequence received packets and to discard duplicates. By 1976, we were working on the second iteration of the design and testing with the Packet Satellite and Packet Radio networks. The latter network was operated by SRI International in the San Francisco Bay area and we used this system to demonstrate mobile access to digital information. By 1977, we implemented the third iteration of the TCP protocol design and separated it into two parts: a simple IP (internet protocol) that delivered packets without any sequencing or attempt to correct for dropped packets and the revised TCP which sat on top of IP and tried to recover from dropped packets and to put packets back in order. This split was important because it allowed for the handling of real-time traffic for which low latency of delivery was more important the delivering precisely and in-order everything sent. From this concept has come support for voice over IP (VOIP) and other real-time applications including streaming video, tracking applications, online games and so on.

At about this time, colleagues at DARPA, including Bob Kahn, began to urge me to come to DARPA to run the Internet research program from there. While reluctant at first, they won me over and I joined DARPA in the late summer of 1976.

On November 22, 1977, we demonstrated all three networks operating together: the ARPANET, the Packet Radio Net and the Packet Satellite Net. This was a significant milestone in the development of the Internet protocols. We continued to test the systems and by mid-1978 we had established the stability of the new TCP/IP protocols.

Then began a heavy effort to implement the protocols on as many operating systems as possible (and there were many in those days). Perhaps among the most important was the UNIX operating system that was developed at AT&T Laboratories. DARPA contracted with the University of California, Berkeley, to implement TCP/IP for Unix and this software was made available, free of charge, for people running workstations using the Unix operating system. I believe that this decision, together with the commercialization of Ethernet, had a strong influence on the adoption of TCP/IP by the academic community.

By 1982, we concluded that there were sufficient interoperable versions of the TCP/IP protocols that we could insist that all the computers on the ARPANET, Packet Radio Net and Packet Satellite Net (and Ethernets) should adopt the new protocols and retire the earlier ARPANET protocol NCP (Network Control Protocol). On January 1, 1983, this cutover was made and although there were some stragglers, virtually all the computers on the network (about

First Three-Network Test of Internet

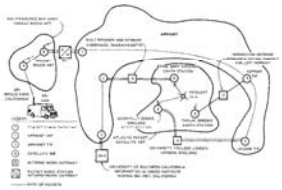

November 22, 1977

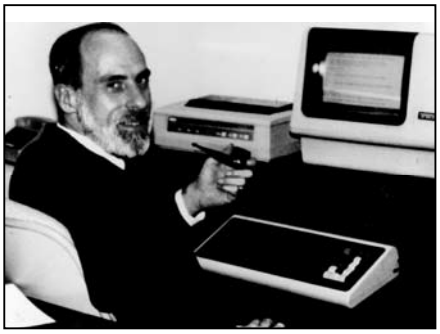

400 of them) were operating TCP/IP by the middle of 1983.

In the meantime, I had moved from DARPA in late 1982 to join MCI where I led the engineering work to design and build a commercial electronic mail system called MCI Mail. This was a rather ambitious and advanced system that allowed for the interconnection of multiple electronic mail systems. Electronic mail messages could be addressed not only to recipients able to receive email but also to Telex users and those who could only receive postal mail. Later the system was enhanced to send facsimile messages as well. By 1986, the system reached a stable operating point and development slowed.

Bob Kahn left DARPA in 1985 and started a new Corporation for National Research Initiatives and invited me to join him there to carry out an ambitious research program on information infrastructure. We explored concepts such as digital libraries, mobile software programs we called Knowledge Robots ("Knowbots"), high speed networking and many other ideas. In this same time period, the US National Science Foundation began the construction of its NSFNET around the TCP/IP protocols. The effort quickly evolved to become a major force in academic networking not only in the United States but also in other countries where academic network might be interconnected with the NSFNET as part of a growing, global Internet. Around 1988, it struck me that the Internet might not become available to the general public without some kind of commercial engine to support it. Believing that it should become a commercially accessible service, I proposed to interconnect the commercial MCI Mail system to the Internet and had to seek permission from the US Government's Federal Networking Council. Until that time, it was not permitted to carry commercial traffic on the NSFNET backbone, so I requested permission to violate that rule as an experiment. By 1989, there were three commercial Internet service providers in operation in the United States and legislation was developed and passed in 1992 to allow the Government-sponsored networks of the US part of the Internet to carry commercial traffic.

It was about this time that Tim Berners-Lee's work on the World Wide Web became visible in the form of the MOSAIC browser version developed by Marc Andreesen and his colleague Eric Bina at the National Center for Supercomputer Applications. Around 1994, Marc joined with James Clark to found Netscape Communications. In the Fall of 1994, I re-joined MCI to help to put the company into the commercial Internet business. With the highly visible Initial Public Offering of Netscape Communications, the "dot-boom" in the United States began. Many Internet-based companies were started and funded, many without much of a business plan to justify the investments.

In this time frame, I received the Kilby award from the inventor of the integrated circuit, Jack Kilby who later went on to receive the Nobel Prize for his work. Bob Kahn and I were honored to receive the US National Medal of Technology from President Bill Clinton in 1997. Many other awards were given to us together and individually during that time. I found my self involved in a science fiction television program developed by Majel Roddenberry, the widow of Gene Roddenberry, the creator of Star Trek and was once again at the White House in 1999 for a Millennium evening hosted by President and Mrs. Clinton.

As the Dot-Boom became the Dot-Bust, many of the Internet-based companies failed but despite this painful outcome, there was a persistent and growing demand for access to the Internet and a creative ferment around the invention of new applications that could take advantage of its growing global connectivity. By 2005, the Internet had become a very visible fixture in the lives of many and for our contributions, Bob and I were once again at the White House receiving the Presidential Medal of Freedom from President George W. Bush.

As the many years have passed the one constant that has remained and influenced my thinking is the long-term collaborative partnership that Bob Kahn and I have enjoyed. If I have learned anything from this many decade long adventure, it is that one's work and success will often rest on relationships established early in life and nourished for a lifetime. The success of the Internet has depended on many such relationships among the millions who contribute to its evolution today. We should value these relationships and encourage them among our colleagues and, perhaps more important, among the younger generations who will succeed us and succeed in their work through their own partnerships and collaborations. I am curious, as I expect you are as well, to see what these young people can accomplish as they set out to explore unknown territory with tools and capabilities we may not even be able to imagine today.